

Distribution Preserving Source Separation with Time Frequency Predictive Models

Pedro J. Villasana T. Janusz Klejsa Lars Villemoes Per Hedelin

Dolby Sweden AB

Stockholm, Sweden

{pedro.villasana; janusz.klejsa; lars.villemoes; per.hedelin}@dolby.com

Abstract—We provide an example of a distribution preserving source separation method, which aims at addressing perceptual shortcomings of state-of-the-art methods. Our approach uses unconditioned generative models of signal sources. Reconstruction is achieved by means of mix-consistent sampling from a distribution conditioned on a realization of a mix. The separated signals follow their respective source distributions, which provides an advantage when separation results are evaluated in a listening test.

Index Terms—generative source separation, generative modeling, Langevin sampling

I. INTRODUCTION

Reconstructions of source signals generated by state-of-the-art audio separation algorithms often sound unnatural. Even a scheme with good objective performance may fall short in a listening test, since its reconstruction may be deemed non-plausible by a human observer. This problem is also discussed in the context of signal enhancement, where a trade-off between distortion and plausibility of the reconstruction is recognized [1]. In this paper, we provide an example of a distribution preserving source separation (DPSS) scheme, where we require that the reconstructed signal components obey their respective source distributions.

Once a mixture of audio signals is created, some information about its components is inevitably lost. Regression-based separation schemes [2] tend to suppress reconstruction when there is an uncertainty about the source. This can, for example, introduce spectral holes, which can be sometimes easily detected while listening, even though the objective performance is high [3].

Generative separation methods were introduced in [4]–[6], where schemes operate on magnitude spectra. Such a setting is generally suboptimal, since it discards the phase information, which may introduce artifacts [7]. A post-processing approach to improve the perceptual performance of spectrogram-based separation was proposed in [8]. In the context of audio separation, generative modeling in the signal domain currently provides the state-of-the-art reconstruction quality [9], [10].

The approach proposed in this paper stems from [9], which uses generative unconditional models of signal components in the construction of a separation scheme. While the separation scenarios of [9], [10] are toy-like (because these approaches assume that generative models of all the mixture components are available), they still provide a clean testing ground for the proposed separation approach. In our case, we

focus on employing high-quality generative models to provide in-distribution reconstructions. Specifically, we use generative models operating in a quadrature modulated filter (QMF) bank domain [11] that were designed to deliver high quality audio synthesis. Our goal is to make the separated components perceptually indistinguishable from observations of signals generated from their sources (as illustrated in Fig. 1). We choose separation of piano-speech mixtures to facilitate comparison with [9], but also to demonstrate that the principle works with diverse signal categories.

There are many objective performance measures for evaluating source separation methods [12]–[14]. However, it is also commonly acknowledged that such measures do not reliably predict the perceptual performance [15]. Here we focus on evaluating the separation results in a listening test. We use the MUSHRA methodology [16], as normally applied for evaluation of audio coding schemes.

The paper is organized as follows. We describe the separation method in Section II. The approach involves usage of source models capable of providing in-distribution reconstructions of the sources. We describe a structure of such a model in Section III. The results of objective evaluation of the scheme and results of the listening tests are provided in Section IV.

II. SEPARATION METHOD

An observation of a downmix $\mathbf{y} = g(\mathbf{x})$ where $g : \mathbb{R}^{S \times N} \rightarrow \mathbb{R}^N$ is the result of a mixing operation of S independent components contained in a realization of a source vector \mathbf{x} . The goal is to reconstruct the components in \mathbf{x} by sampling from a trained model $p(\mathbf{x}|\mathbf{y})$. We follow the approach of [9], where Langevin sampling facilitates working with the log gradient of the posterior of the joint distribution:

$$\nabla_{\mathbf{x}} \log p(\mathbf{x}, \mathbf{y}) = \nabla_{\mathbf{x}} \log p(\mathbf{y}|\mathbf{x}) + \nabla_{\mathbf{x}} \log p(\mathbf{x}). \quad (1)$$

Such a sampling scheme can be implemented by using a score matching model [10], [17], [18], where $\nabla_{\mathbf{x}} \log p(\mathbf{x})$ would be the score models for the components of the mix (which would be the only trainable parts of the scheme). Alternatively, the gradients can be computed in runtime by using trained generative models $p(\mathbf{x})$.




Fig. 1. Spectrograms: (a) Input mixture (speech+piano); (b) Separated piano (DPSS); (c) Separated speech (DPSS).

A. Langevin sampling

Due to (1), the Langevin sampling for sufficiently small η can be implemented as:

$$\begin{aligned} \underline{\mathbf{x}}_{t_i} &= \underline{\mathbf{x}}_{t_{i-1}} + \eta \nabla_{\underline{\mathbf{x}}} \log p(\underline{\mathbf{x}}_{t_{i-1}} | \mathbf{y}) + \sqrt{2\eta} \underline{\boldsymbol{\epsilon}}_{t_i} \\ &= \underline{\mathbf{x}}_{t_{i-1}} + \eta \nabla_{\underline{\mathbf{x}}} \left(\log p(\underline{\mathbf{x}}_{t_{i-1}}) + \log p(\mathbf{y} | \underline{\mathbf{x}}_{t_{i-1}}) \right) + \sqrt{2\eta} \underline{\boldsymbol{\epsilon}}_{t_i} \end{aligned} \quad (2)$$

where $i = 1, \dots, I$ is the iteration number, $t \in \mathcal{U}(0, 1)$ describes the algorithmic time going from $t_0 = 0$ to $t_I = 1$, and $\underline{\boldsymbol{\epsilon}} \in \mathbb{R}^{S \times N}$ is a vector of S noise realizations sampled from $\mathcal{N}(\mathbf{0}, \mathbf{I})$. The scheme of [9] used auto-regressive models to describe $p(\underline{\mathbf{x}})$. However, in that case $\underline{\mathbf{x}}$ was discrete and therefore non-differentiable. This problem was alleviated by sampling from a smoothed version of $\underline{\mathbf{x}}$, where $\tilde{\underline{\mathbf{x}}} = \underline{\mathbf{x}} + \sigma_m \underline{\mathbf{z}}$ and $\underline{\mathbf{z}} \in \mathbb{R}^{S \times N}$ consisting of S noise realizations sampled from $\mathcal{N}(\mathbf{0}, \mathbf{I})$, for $\sigma_1 > \dots > \sigma_M$. This requires using a collection of M models for each source. Here, instead, we use noisy-predictors described in Section II-C, where there is only a single model per source.

B. Auto-regressive source model

While the separation approach is compatible with modern score matching models such as [18], [10], here, we use auto-regressive (AR) modeling of signals (as in [9]), since they seem to have good enough performance for our task. However, our AR models operate in a filter-bank domain on a continuous signal representation. It is known that time-domain AR modeling of signals that are close to being a sum of AR processes (e.g., piano) is not efficient [19]. The decorrelating property of the filter-bank seems to aid the model efficiency for such complex sources in our experiments.

In our model, a probability distribution of a random variable represented by a matrix of time-frequency tiles of a signal representation provided by the filter-bank is given by a product of probabilities of conditionally independent frames:

$$p(\mathbf{X}) = \prod_n p(\mathbf{x}_n | \mathbf{X}_{<n}), \quad (3)$$

where \mathbf{x}_n denotes the frequency coefficients at time frame n and $\mathbf{X}_{<n}$ comprises past time frames. This model structure resembles an MDCT model of [20], but in this work we use

another implementation designed for a real-valued QMF. We discuss implementation of the model in Section III.

C. Noisy predictor

The Langevin sampling method depends on providing a description of $p(\mathbf{x}_{t_i})$. This means that we must have models that are able to describe the probability of \mathbf{x} in between \mathbf{x}_0 (noise) and \mathbf{x}_1 (source). Here we perform AR modelling for noisy sources $\tilde{\mathbf{x}} = \mathbf{x} + \sigma \mathbf{z}$ with $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ conditioned on the noise level, i.e. $p(\tilde{\mathbf{x}} | \sigma)$. This helps to have more control over the Langevin sampling algorithm, thus modifying (2) to

$$\begin{aligned} \tilde{\underline{\mathbf{x}}}_{t_i} &= \tilde{\underline{\mathbf{x}}}_{t_{i-1}} + \eta \nabla_{\tilde{\underline{\mathbf{x}}}} \log p(\tilde{\underline{\mathbf{x}}}_{t_{i-1}} | \mathbf{y}, \sigma_{t_i}) + \sqrt{2\eta} \underline{\boldsymbol{\epsilon}}_{t_i} \\ &= \tilde{\underline{\mathbf{x}}}_{t_{i-1}} + \eta \nabla_{\tilde{\underline{\mathbf{x}}}} \left(\log p(\tilde{\underline{\mathbf{x}}}_{t_{i-1}} | \sigma_{t_i}) + \log p(\mathbf{y} | \tilde{\underline{\mathbf{x}}}_{t_i}, \sigma_{t_i}) \right) \\ &\quad + \sqrt{2\eta} \underline{\boldsymbol{\epsilon}}_{t_i}. \end{aligned} \quad (4)$$

Writing the linear function g as $g(\underline{\mathbf{x}}) = \underline{\mathbf{a}}^T \underline{\mathbf{x}}$ we have that

$$\mathbf{y} = g(\underline{\mathbf{x}}) = g(\tilde{\underline{\mathbf{x}}} - \sigma \underline{\mathbf{z}}) \sim \mathcal{N}(g(\tilde{\underline{\mathbf{x}}}), \sigma^2 \|\underline{\mathbf{a}}\|^2 \mathbf{I}), \quad (5)$$

where $\underline{\mathbf{a}} \in \mathbb{R}^S$ defines the weights for mixing the sources. The proposed construction eliminates the need for training a sequence of de-noising models as in [9].

D. Scheduler

Usually the steps $\sigma_{t_0} > \sigma_{t_1} > \dots > \sigma_{t_I}$ in (4) are geometrically spaced by $\gamma = \sigma_{t_{i+1}} / \sigma_{t_i}$ with $t_0 = 0$ and $t_I = 1$. Typically, (4) is applied using annealed Langevin sampling (ALS) [21] with $\eta = \varepsilon_\eta \sigma_{t_i} / \sigma_{t_I} = \varepsilon_\eta \gamma^{i-I}$ where ε_η is a hyperparameter.

In our scheme we perform Langevin sampling through consistent annealed sampling (CAS):

$$\begin{aligned} \tilde{\underline{\mathbf{x}}}_{t_i} &= \tilde{\underline{\mathbf{x}}}_{t_{i-1}} + \alpha \sigma_{t_i}^2 \nabla_{\tilde{\underline{\mathbf{x}}}} \log p(\tilde{\underline{\mathbf{x}}}_{t_{i-1}} | \mathbf{y}, \sigma_{t_i}) + \beta \sigma_{t_{i+1}} \underline{\boldsymbol{\epsilon}}_{t_i} \\ &= \tilde{\underline{\mathbf{x}}}_{t_{i-1}} + \beta \sigma_{t_{i+1}} \underline{\boldsymbol{\epsilon}}_{t_i} \\ &\quad + \alpha \sigma_{t_i}^2 \nabla_{\tilde{\underline{\mathbf{x}}}} \left(\log p(\tilde{\underline{\mathbf{x}}}_{t_{i-1}} | \sigma_{t_i}) + \log p(\mathbf{y} | \tilde{\underline{\mathbf{x}}}_{t_{i-1}}, \sigma_{t_i}) \right), \end{aligned} \quad (6)$$

where α and β can be re-parameterized as shown in [22] by $\alpha = 1 - \gamma^\eta$, $\beta = \sqrt{1 - \gamma^{2(\eta-1)}}$ with $\eta \geq 1$ for $i \in \{1, 2, \dots, I-1\}$, and $\alpha = 1$, $\beta = 0$ for $i = I$ ($t = 1$).

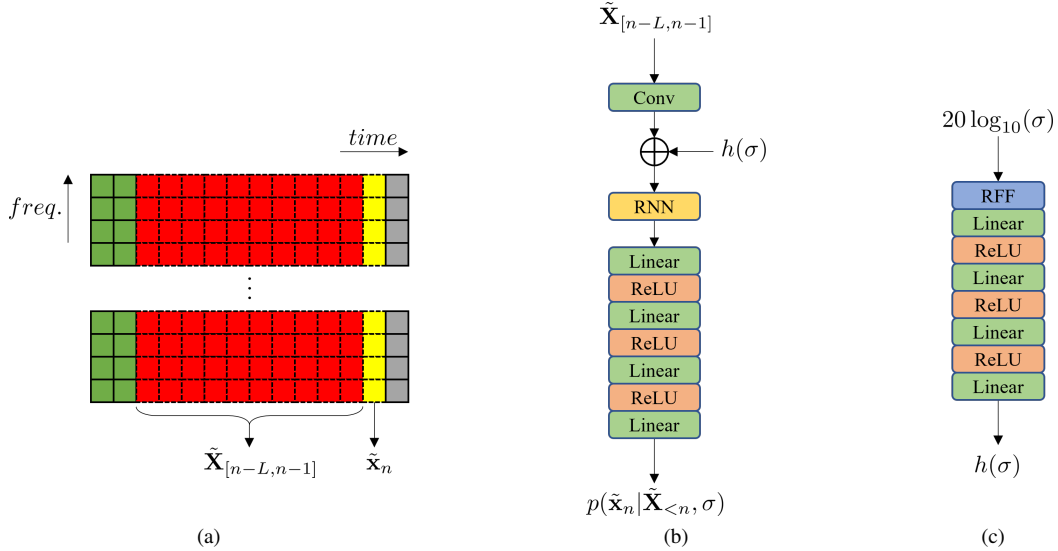


Fig. 2. Source model operation: (a) Generation example; (b) Prediction network; (c) Conditioning network

III. SOURCE MODEL

We designed an auto-regressive model that operates in the QMF domain, predicting noisy frames conditioned on the noise level, i.e. the model performs

$$p(\tilde{\mathbf{X}}|\sigma) = \prod_n p(\tilde{\mathbf{x}}_n|\tilde{\mathbf{X}}_{<n}, \sigma) \quad (7)$$

with $\tilde{\mathbf{X}}$ being the result of adding Gaussian noise of power σ^2 to the clean signal \mathbf{X} . This is a simplified version of the model in [20] where they do frequency prediction within a time frame. Here, we predict the whole frame in one shot.

A high-level view on the model architecture is shown in Fig. 2. Here, we highlight the two aspects of the model: its predictive structure and noise-level conditioning.

Prediction. Most natural occurring audio signals have certain correlation in time and frequency, examples of these are the tempo of a song and the timbre of a musical instrument, respectively. To exploit this correlation we designed the model to perform prediction of vectors. Looking at Fig. 2a, to predict the samples $\tilde{\mathbf{x}}_n$ (yellow) in frame n , the model takes L frames in the past $\tilde{\mathbf{X}}_{[n-L, n-1]}$ (red). These samples are passed through a convolutional layer and the output is added to the output of the conditioning network. The resulting sum goes into an RNN that handles coherence over time. Then the output of the RNN is put through an MLP layer consisting of 4 linear layers with ReLUs in between them. The output of the final linear layer contains the parameters necessary to create the multivariate distribution $p(\tilde{\mathbf{x}}_n|\tilde{\mathbf{X}}_{<n}, \sigma)$. The prediction network is shown in Fig. 2b.

Noise-level conditioning. The conditioning network shown in Fig. 2c takes the noise level σ in dB and converts it into Random Fourier Features (RFF) [23]. These features are then passed through a 4-layer MLP with ReLUs in between each linear layer.

IV. EVALUATION

We evaluated the proposed method on mixtures created from piano and speech signals. The piano signals came from the Supra database [24], while the speech excerpts came from the VCTK database [25]. This selection allows us to demonstrate the gain in objective and subjective performance, when the separation framework of [9] is configured towards source distribution preservation by employing high-quality source models.

A. Model setup and training

We used a QMF bank of 64 channels. The model was designed to predict 1 frame at a time by looking at 10 frames in the past ($L = 10$ in Fig. 2a). We opted for employing a single-layer LSTM as RNN. All elements in the model use a hidden dimension of size 1024 and the output stage of the last linear layer in Fig. 2b produced 128 parameters. These parameters represent means μ and scales s used to construct Logistic distributions $\text{Log}(x; \mu, s) = \frac{1}{4s} \text{sech}^2\left(\frac{x-\mu}{2s}\right)$ for each of the 64 QMF channels in the frame. We opted for using the Logistic distribution due to its smooth and range-limited derivative $\frac{\partial}{\partial x} \log \text{Log}(x; \mu, s) = -\frac{1}{s} \tanh\left(\frac{x-\mu}{2s}\right)$. This property proved beneficial when computing $\nabla_{\tilde{\mathbf{X}}} \log p(\tilde{\mathbf{X}}|\sigma)$ for the application of (6).

We trained source models using the VCTK speech and Supra piano datasets. Each model had approximately 17M parameters. The datasets were randomly divided into training, validation and test sets, consisting of 80-10-10 percent of the data, respectively. All items were adjusted to a length of 8 seconds by concatenation or truncation. We trained each model using NLL loss for 1M iterations, using the ADAM optimizer [26], starting with a learning rate of 1e-4 and reducing it with a cosine scheduler [27] to a final value of 1e-6. Each iteration consists of a batch of 64 items, using a sequence length of 1 second. Every file in the batch was corrupted by adding

TABLE I
SI-SDR [dB] FOR IRM, PNF AND THE PROPOSED METHOD (DPSS)

	IRM	PNF [9]	DPSS
speech	16.25	15.93	22.43
piano	13.32	12.87	19.59
mix	64.52	38.92	63.34

Gaussian noise with power distributed uniformly in the range [-90, 0] dB. We kept this noise level for the entire duration of the file, such that 8 consecutive iterations belonged to the same items and noise levels.

When performing source separation, we used the CAS scheduler (6) starting from $\sigma_0 = 0$ dB to $\sigma_1 = -90$ dB, using $\eta = 90$ for $I = 1500$ iterations. We noted that the results varied significantly depending on the power level of the individual sources. Therefore, we decided to normalize the mix to a level of -23 dB power before applying (6), and de-normalize the estimated sources afterwards.

B. Objective performance

We evaluated the separation performance by using SI-SDR [14] as an objective measure. We constructed a regression-based benchmark with an ideal ratio mask (IRM) computed on per-bin basis with a 1024 samples stride. We used a time-domain scheme of [9] (PNF), which is based on WaveNet [28] operating with a discrete output stage (8-bit PCM), where we computed the SI-SDR with respect to the original signal (16-bit PCM). Finally, we included the proposed scheme. Results of these three approaches are shown in Table I for speech, piano, and the mix. It can be seen that the proposed scheme has a superior performance for both source categories. But as expected, an ideal regression scheme has the best mix consistency. The mix consistency of the PNF condition is limited by the discrete output stage of its generative models.

C. Listening test

We analyzed the performance of the separation methods by two MUSHRA [16] tests. We created a set of test mixture signals by using random items from test sets of the respective source models.¹ The listeners listened separately to separated components for each of the sources in separate sessions, and they were not exposed to the mixture signals. Hence, the setup resembled the design of listening tests commonly used for evaluation of audio codecs. While this way of testing is challenging for separation algorithms, it allows us to test the hypothesis that the proposed approach can operate in the high quality regime.

The conditions in the tests included: a hidden reference, a 3.5 kHz low-pass anchor (LP35), and the signals extracted by the three separation methods. The anchor signals were created from the references (not from the mixture signals). 11 expert listeners participated in both listening tests.

¹Excerpts from the test are available here: <https://dpss-demo.github.io/>

The results of the listening test for separated piano are shown in Fig. 3 and the ones for separated speech are shown in Fig. 4. For both source categories, it can be seen that DPSS provides a significant perceptual advantage, and outperforms the benchmarks. The unusually high score of the low pass anchor in the piano test is due to fact that piano signals in the Supra database have band-limited characteristics.

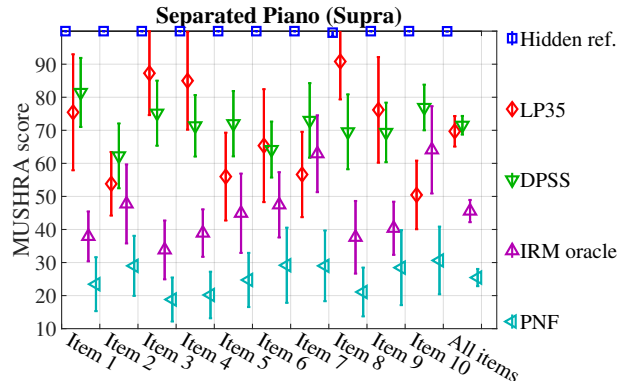


Fig. 3. Listening test results for Supra piano (11 listeners, 95% confidence intervals, Student's t-distribution).

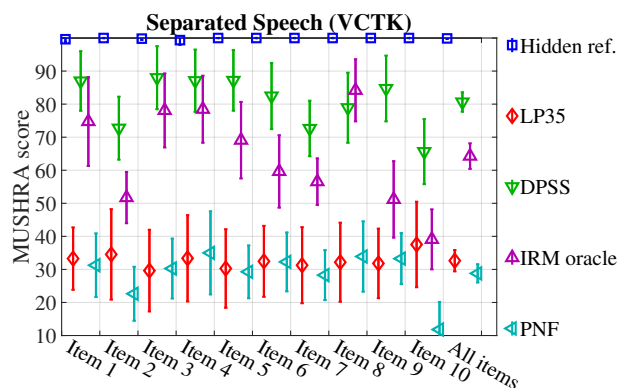


Fig. 4. Listening test results for VCTK speech (11 listeners, 95% confidence intervals, Student's t-distribution).

V. CONCLUSION

We conclude that preservation of source distributions can be a viable strategy for separation of audio mixtures, as it facilitates generation of plausible reconstructions of the sources. We demonstrate its advantage by performing MUSHRA tests in a rigorous setup commonly used for evaluation of audio codecs. Furthermore, our results indicate that the proposed filter-bank domain model with a continuous output stage can significantly outperform the time-domain benchmark with a discrete output stage.

REFERENCES

- [1] Y. Blau and T. Michaeli, "The perception-distortion tradeoff," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 6228–6237.

- [2] E. Gusó, J. Pons, S. Pascual, and J. Serrà, "On loss functions and evaluation metrics for music source separation," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 306–310.
- [3] Y. Luo and N. Mesgarani, "TaSNet: Time-domain audio separation network for real-time, single-channel speech separation," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 696–700.
- [4] Y. C. Subakan and P. Smaragdis, "Generative adversarial source separation," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 26–30.
- [5] V. Narayanaswamy, J. J. Thiagarajan, R. Anirudh, and A. Spanias, "Unsupervised audio source separation using generative priors," *arXiv preprint arXiv:2005.13769*, 2020.
- [6] C. Hawthorne, I. Simon, A. Roberts, N. Zeghidour, J. Gardner, E. Manilow, and J. Engel, "Multi-instrument music synthesis with spectrogram diffusion," *arXiv preprint arXiv:2206.05408*, 2022.
- [7] F. Lluís, J. Pons, and X. Serra, "End-to-end music source separation: is it possible in the waveform domain?" *arXiv preprint arXiv:1810.12187*, 2018.
- [8] N. Schaffer, B. Cogan, E. Manilow, M. Morrison, P. Seetharaman, and B. Pardo, "Music Separation Enhancement with Generative Modeling," *arXiv preprint arXiv:2208.12387*, 2022.
- [9] V. Jayaram and J. Thickstun, "Parallel and flexible sampling from autoregressive models via Langevin dynamics," in *International Conference on Machine Learning*. PMLR, 2021, pp. 4807–4818.
- [10] G. Mariani, I. Tallini, E. Postolache, M. Mancusi, L. Cosmo, and E. Rodolà, "Multi-source diffusion models for simultaneous music generation and separation," *arXiv preprint arXiv:2302.02257*, 2023.
- [11] P. P. Vaidyanathan, *Multirate Systems and Filter Banks*. Prentice Hall: Englewood Cliffs, NJ, 1994.
- [12] E. Cano, D. Fitzgerald, and K. Brandenburg, "Evaluation of quality of sound source separation algorithms: Human perception vs quantitative metrics," in *2016 24th European Signal Processing Conference (EUSIPCO)*. IEEE, 2016, pp. 1758–1762.
- [13] T. Kastner and J. Herre, "An efficient model for estimating subjective quality of separated audio source signals," in *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2019, pp. 95–99.
- [14] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "SDR—half-baked or well done?" in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 626–630.
- [15] M. Torcoli, T. Kastner, and J. Herre, "Objective measures of perceptual audio quality reviewed: An evaluation of their application domain dependence," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1530–1541, 2021.
- [16] "Method for the subjective assessment of intermediate quality levels of coding systems," Recommendation ITU-R BS.1534-3, Oct. 2015.
- [17] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," *arXiv preprint arXiv:2011.13456*, 2020.
- [18] S. Pascual, G. Bhattacharya, C. Yeh, J. Pons, and J. Serrà, "Full-band general audio synthesis with score-based diffusion," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [19] C. W. Granger and M. J. Morris, "Time series modelling and interpretation," *Journal of the Royal Statistical Society: Series A (General)*, vol. 139, no. 2, pp. 246–257, 1976.
- [20] G. Davidson, M. Vinton, P. Ekstrand, C. Zhou, L. Villemoes, and L. Lu, "High quality audio coding with MDCTNet," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [21] Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," *Advances in neural information processing systems*, vol. 32, 2019.
- [22] J. Serrà, S. Pascual, and J. Pons, "On tuning consistent annealed sampling for denoising score matching," *arXiv preprint arXiv:2104.03725*, 2021.
- [23] S. Rouard and G. Hadjeres, "CRASH: raw audio score-based generative modeling for controllable high-resolution drum sound synthesis," *CoRR*, vol. abs/2106.07431, 2021. [Online]. Available: <https://arxiv.org/abs/2106.07431>
- [24] Z. Shi, C. Sapp, K. Arul, J. McBride, and J. O. Smith III, "SUPRA: Digitizing the Stanford University Piano Roll Archive," in *ISMIR*, 2019, pp. 517–523.
- [25] J. Yamagishi, C. Veaux, K. MacDonald *et al.*, "CSTR VCTK Corpus: English multi-speaker corpus for CSTR voice cloning toolkit (version 0.92)," *University of Edinburgh. The Centre for Speech Technology Research (CSTR)*, 2019.
- [26] D. P. Kingma and J. Ba, "ADAM: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [27] I. Loshchilov and F. Hutter, "SGDR: Stochastic gradient descent with warm restarts," *arXiv preprint arXiv:1608.03983*, 2016.
- [28] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, "WaveNet: A generative model for raw audio," *CoRR*, vol. abs/1609.03499, 2016. [Online]. Available: <http://arxiv.org/abs/1609.03499>